A Theoretical Analysis of First Heuristics of Crowdsourced Entity Resolution

Arya Mazumdar and **Barna Saha**
College of Information & Computer Sciences
University of Massachusetts Amherst
{arya,barna}@cs.umass.edu

arXiv:1702.01208v1 [cs.DB] 3 Feb 2017

Abstract

Entity resolution (ER) is the task of identifying all records in a database that refer to the same underlying entity, and are therefore duplicates of each other. Due to inherent ambiguity of data representation and poor data quality, ER is a challenging task for any automated process. As a remedy, human-powered ER via crowdsourcing has become popular in recent years. Using crowd to answer queries is costly and time consuming. Furthermore, crowd-answers can often be faulty. Therefore, crowd-based ER methods aim to minimize human participation without sacrificing the quality and use a computer generated similarity matrix actively. While, some of these methods perform well in practice, no theoretical analysis exists for them, and further their worst case performances do not reflect the experimental findings. This creates a disparity in the understanding of the popular heuristics for this problem. In this paper, we make the first attempt to close this gap. We provide a thorough analysis of the prominent heuristic algorithms for crowd-based ER. We justify experimental observations with our analysis and information theoretic lower bounds.

1 Introduction

Entity resolution (ER, record linkage, deduplication, etc.) seeks to identify which records in a data set refer to the same underlying real-world entity (Fellegi and Sunter 1969; Elmagarmid, Ipeirotis, and Verykios 2007; Getoor and Machanavajjhala 2012; Larsen and Rubin 2001; Christen 2012). Our ability to represent information about real-world entities in very diverse ways makes this a complicated problem. For example, collecting profiles of people and businesses, or specifications of products and services from websites and social media sites can result in billions of records that need to be resolved. These entities are identified in a wide variety of ways, complicated further by language ambiguity, poor data entry, missing values, changing attributes and formatting issues. ER is a fundamental task in data processing with wide-array of applications. There is a huge literature on ER techniques; many include machine learning algorithms, such as decision trees, SVMs, ensembles of classifiers, conditional random fields, unsupervised learning etc. (see (Getoor and Machanavajjhala 2012) for a recent survey). Yet, ER remains a demanding task for any automated strategy yielding low accuracy.

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ER can be cast as a clustering problem. Consider a set of n elements V that must be clustered into k disjoint parts $V_i, i = 1, 2, \ldots, k$. The true underlying clusters $V_i \in [n]$, $i \in [1, k]$ are unknown to us, and so is k. Each of these V_is represents an entity. Each element $v \in V$ has a set of attributes. A similarity function is used to estimate the similarity of the attribute sets of two nodes u and v. If u and v represent the same entity, then an ideal similarity function will return 1, and if they are different, then it will return 0. However, in practice, it is impossible to find an ideal similarity function, or even a function close to it. Often, some attribute values may be missing or incorrect, and that leads to similarity values that are noisy representation of the ideal similarity function. Any automated process that uses such similarity function is thus prone to make errors. To overcome this difficulty, a relatively recent line of works propose to use human knowledge via crowdsourcing to boost accuracy of ER (Davidson et al. 2014; Firmani, Saha, and Srivastava 2016; Verroios and Garcia-Molina 2015; Gruenheid et al. 2015; Wang et al. 2012; Wang et al. 2013; Vesdapunt, Bellare, and Dalvi 2014; Yi et al. 2012; Whang, Lofgren, and Garcia-Molina 2013). Human based on domain knowledge can match and distinguish entities with complex representations, where automated strategies fail.

Motivating example. Consider the following illustrative example shown in Figure 1. The Walt Disney, commonly known as Disney, is an American multinational media and entertainment company that owns and licenses 14 theme parks around the world. [1] Given the six places (r_1) Disney World, (r_2) Walt Disney World Resort, (r_3) Walt Disney Theme Park, Orlando, (r_4) Disneyland, (r_5) Disneyland Park, humans can determine using domain knowledge that these correspond to two entities: r_1, r_2, r_3 refer to one entity, and r_4, r_5 refer to a second entity.

Answering queries by crowd could be time-consuming and costly. Therefore, a crowd based ER strategy must attempt to minimize the number of queries to the oracle while resolving the clusters exactly. Having access to ideal crowd answers, a good ordering of comparing record pairs is (r_1, r_2), (r_2, r_3), (r_4, r_5), (r_1, r_5). After the first three pairs have been compared, we can safely infer as "matching" the remaining pair

[1] https://en.wikipedia.org/wiki/The_Walt_Disney_Company

Figure 1: Record pairs connected by green (resp. red) edges are "matching" (resp. "non-matching") in the real world. The numbers on the edges indicate the ordering of queries. While querying strategy of (a) only results in 4 queries, the querying strategy of (b) results in 9 queries.

(r_1, r_3) leveraging transitive relations. After the last pair in the ordering has been compared, we can safely infer as "non-matching" all the remaining pairs (r_1, r_4), (r_2, r_4), (r_2, r_5), (r_3, r_4), (r_3, r_5) in the database.

The work by Wang et al. (Wang et al. 2013) was among the first few (Wang et al. 2012; Demartini, Difallah, and Cudré-Mauroux 2012; Whang, Lofgren, and Garcia-Molina 2013) to propose the notion of hybrid human-machine approach for entity resolution. Moreover, it is the first paper to leverage the transitive relationship among the entities to minimize the number of queries which has since become a staple in every follow-up work on this topic (Firmani, Saha, and Srivastava 2016; Verroios and Garcia-Molina 2015; Gruenheid et al. 2015; Vesdapunt, Bellare, and Dalvi 2014). Assuming there is an oracle, an abstraction of a crowd-sourcing platform that can correctly answer questions of the form *"Do records u and v refer to the same entity?"*, they presented a new algorithm for crowd-sourced ER. To minimize the number of queries to the crowd oracle, Wang et al. utilizes the transitive relation in which known match and non-match labels on some record pairs can be used to automatically infer match or non-match labels on other record pairs. In short, the heuristic algorithm by Wang et al. does the following: it orders the tuples (record pairs/edges) in nonincreasing order of similarity, and query any edge according to that order whenever the right value of that edge cannot be transitively deduced from the already queried/inferred edges so far.

While the crowd-sourcing algorithm of Wang et al. works reasonably well on real datasets, theoretical guarantees for it was not provided. However, in (Vesdapunt, Bellare, and Dalvi 2014), Vesdapunt et al. showed that in some instances this algorithm can only give an $\Theta(n)$ approximation, that is when an optimum algorithm may require c queries, Wang et al.'s algorithm can require $\Theta(cn)$ queries.

Vesdapunt et al. proposed an algorithm that proceeds in the following iterative manner. In each round, an element to be clustered is compared with one representative of all the existing clusters. The order of these comparisons is defined by a descending order of the similarity measures. As soon as a positive query result is found the element is assigned to the corresponding cluster and the algorithm moves to the next round with a new element. It is easy to see that in the worst case the number of queries made by the algorithm is nk, where n is the number of elements and k is the number of clusters. It also follows that this is at least an $O(k)$ approximation.

Note that (Wang et al. 2013; Vesdapunt, Bellare, and Dalvi 2014) consider the answers of queries are correct as an ideal crowd abstraction - this can often be guaranteed via majority voting. But it is unclear that how the quality of the similarity measurements affects the total number of queries. Indeed, in typical datasets, the performances of the algorithms of Wang et al. and Vesdapunt et al. are quite similar, and they are much better than their worst case guarantees that do not take into account the existence of any meaningful similarity measures. This means the presence of the similarity measures helps reduce the query complexity significantly. Is there a way to theoretically establish that and come up with guarantees that match the experimental observations?

It is of paramount interest to characterize the query complexity (number of questions asked to the crowd) of these popular heuristics and come up with algorithms that minimize such complexity. The query complexity is directly proportional to the overall cost of a crowd-based algorithm, due to the fact that crowd questions are time-consuming and in many times involve compensations. Designing a strategy that would minimize the query complexity can directly be seen as alternatives to *active learning* problem with minimum labeling requirements (Sarawagi and Bhamidipaty 2002; Bellare et al. 2012). From the perspective of lower bounding the query complexity, ER can be seen as a *reinforcement learning* problem. Indeed, in each step of assigning a record to one of the underlying entities, a query must be made wisely so that under any adversarial configurations, the total number of queries remain small.

Contributions. In this paper we assume the following model for the similarity measurements. Let $W = \{w_{u,v}\}_{(u,v)\in V\times V}$ denote the matrix obtained by pair-wise similarity computation, where $w_{u,v}$ is a random variable drawn from a probability distribution f_g if u and v belong to the same cluster and drawn from a probability distribution f_r otherwise. The subscripts of f_r and f_g are chosen to respectively signify a "red edge" (or absence of a link) and a "green edge" (or presence of a link). Note that, this model of similarity matrix is by no means the only possible; however it captures the essential flavor of the problem.

Our main contribution in this paper is to provide a theoretical analysis of query complexities of the two aforementioned heuristics from (Wang et al. 2013; Vesdapunt, Bellare, and Dalvi 2014). Our analysis quantifies the effect of the presence of similarity measures in these algorithms, establishes the superiority between these algorithms under different criteria, and derives the exact expression of query complexity under some fundamental probability models.

Next, to establish the near-optimality or sub-optimality of the above heuristics, we compare our results with an information theoretic lower bound recently proposed by us (Mazumdar and Saha 2016). As a corollary to the results of (Mazumdar and Saha 2016), it can be seen that the information theoretic lower bound depends on the Hellinger divergence between f_g and f_r. More interestingly, the quality of the similarity matrix can be characterized by the Hellinger divergence between f_g and f_r as well.

Finally, we show that the experimental observations of (Wang et al. 2013; Vesdapunt, Bellare, and Dalvi 2014) agree with our theoretical analysis of their algorithms. Moreover, we conduct a thorough experiment on the bibliographical `cora` (McCallum 2004) dataset for ER and several synthetic datasets to validate the theoretical findings further.

2 System model and techniques

2.1 Crowdsourced Entity Resolution Crowd-ER

Consider a set of elements $V \equiv [n]$ which is a disjoint union of k clusters V_i, $i = 1, \ldots, k$, where k and the subsets $V_i \subseteq [n]$ are unknown. The crowd (or the oracle) is presented with an element-pair $(u, v) \in V \times V$ for a query, that results in a binary answer denoting the event u, v belonging to the same cluster. Note that, this perfect oracle model had been used in the prominent previous works by Wang et al. and Vesdapunt et al.(2013; 2014).

We can assume that with probability $0 < p_i < 1$, the crowd gives a wrong answer to the ith query. However, with resampling the ith query $\Omega(\log n)$ times, that is by asking the same ith query to $\Omega(\log n)$ different users and by taking the majority vote, we can drive the probability p_i to nearly 0 and return to the model of perfect oracle. Note that we have assumed independence among the resampled queries over the index j, which can be justified since we are sampling a growing ($\Omega(\log n)$) number of samples. Furthermore, repetition of the same query to the crowd may not not lead to reduction in the error probability, i.e., a persistent error. Even in this scenario an element can be queried with multiple elements from a same cluster to infer with certainty whether the element belong to the cluster or not. These situations have been covered in detail in our recent work (Mazumdar and Saha 2016). Henceforth, in this paper, we only consider the perfect oracle model. All our results hold for the faulty oracle model described above with only an $O(\log n)$ blow-up in the query complexity.

Consider W, an $n \times n$ similarity matrix, with the (u, v)th entry $w_{u,v}$ a nonnegative random variable in $[0, 1]$ drawn from a probability density or mass function f_g when u, v belong to the same cluster, and drawn from a probability density or mass function f_r otherwise. f_g and f_r are unknown.

The problem of Crowd-ER is to design a set of queries in $V \times V$, given V and W, such that from the answers to the queries, it is possible to recover V_i, $i = 1, 2, ..., k$.

2.2 The two heuristic algorithms

The Edge ordering algorithm (Wang et al. 2013). In this algorithm, we arrange the set $V \times V$ in non-increasing order of similarity values $w_{i,j}$s. We then query sequentially according to this order. Whenever possible we apply transitive relation to infer edges. For example, if the queries (i, j) and (j, l) both get positive answers then there must be an edge (i, l), and we do not have to make the query (i, l). We stop when all the edges are either queried, or inferred.

The Node ordering algorithm (Vesdapunt, Bellare, and Dalvi 2014). In this algorithm, the empirical expected size of the cluster containing element i, $1 \leq i \leq n$, is first computed as $\sum_j w_{i,j}$. Then all the elements are ordered non-increasingly according to the empirical expected sizes of the clusters containing them. At any point in the execution, the algorithm maintains at most k clusters. The algorithm selects the next element and issues queries involving that element and elements which are already clustered in non-increasing order of their similarity, and apply transitivity for inference. Therefore, the algorithm issues at most one query involving the current node and an existing cluster. Trivially, this gives an $O(k)$-approximation.

3 Information theoretic lower bound

Note that, in the absence of similarity matrix W, any optimal (possibly randomized) algorithm must make $\Omega(nk)$ queries to solve Crowd-ER. This is true because an input can always be generated that makes $\Omega(n)$ vertices to be involved in $\Omega(k)$ queries before they can be correctly assigned. However, when we are allowed to use the similarity matrix, this bound can be significantly reduced. Indeed, the following lower bound follows as a corollary of the results of our previous work (Mazumdar and Saha 2016).

Theorem 1. *Given the number of clusters k and f_g, f_r, any randomized algorithm that does not perform at least $\Omega\Big(\min\Big\{\frac{k^2}{\mathcal{H}^2(f_g,f_r)}, nk\Big\}\Big)$ queries, will be unable to return the correct clustering with high probability, where $\mathcal{H}^2(f_g, f_r) \equiv \frac{1}{2}\int_{-\infty}^{\infty}(\sqrt{f_g(x)} - \sqrt{f_r(x)})^2 dx$ is the squared Hellinger divergence between the probability measures f_g and f_r.*

The main idea of proving this lower bound already appears in our recent work (Mazumdar and Saha 2016), and we give a brief sketch of the proof below for the interested readers. Strikingly, Hellinger divergence between f_g and f_r appears to be the right distinguishing measure even for analyzing the heuristic algorithms.

To show the lower bound we consider an input where one of the clusters are fully formed and given to us. The remaining $k-1$ clusters each has size $a = \Big\lfloor \frac{1}{8\mathcal{H}^2(f_g,f_r)} \Big\rfloor$. We prove the result through contradiction. Assume there exists a randomized algorithm ALG that makes a total of $o\Big(\frac{k^2}{\mathcal{H}(f_g,f_r)^2}\Big)$

queries and assigns all the remaining vertices to correct clusters with high probability. However, that implies that the average number of queries ALG makes to assign each of the remaining elements to a cluster must be $o(k)$.

Since there are k clusters, this actually guarantees the existence of an element that is not queried with the correct cluster V_i it is from, and that completely relies on the W matrix for the correct assignment. Now the probability distribution (which is a product measure) of W, P_W, can be one of two different distributions, P'_W and P''_W depending on whether this vertex belong to V_i or not. Therefore these two distributions must be far apart in terms of total variation distance for correct assignment.

However, the total variation distance between P'_W and P''_W $\|P'_W - P''_W\|_{TV} \leq \sqrt{2}\mathcal{H}(P'_W, P''_W)$. But as both P'_W, P''_W are product measures that can differ in at most $2a$ random variables (recall the clusters are all of size a), we must have, using the properties of the Hellinger divergence, $\mathcal{H}(P'_W, P''_W) \leq \sqrt{2a\mathcal{H}(f_g, f_r)^2} \leq \frac{1}{2}$. This means, $\|P'_W - P''_W\|_{TV} \leq \frac{1}{\sqrt{2}}$, i.e., the two distributions are close enough to be confused with a positive probability - which leads to a contradiction. Note that, in stead of recovery with positive probability, if we want to ensure exact recovery of the clusters (i.e., with probability 1) we must query each element at least once. This leads to the following corollary.

Corollary 1. *Any (possibly randomized) algorithm with the knowledge of f_g, f_r, and the number of clusters k, must perform at least $\Omega\Big(n + \frac{k^2}{\mathcal{H}^2(f_g,f_r)}\Big)$ queries, $\mathcal{H}(f_g, f_r) > 0$, to return the correct clustering exactly.*

4 Main results: Analysis of the heuristics

We provide expressions for query complexities for both the edge ordering and the node ordering algorithms. It turns out that the following quantity plays a crucial role in the analysis of both:

$$L_{g,r}(t) \equiv \int_0^1 \Big(\int_0^r f_g(y)dy\Big)^t f_r(x)dx.$$

Theorem 2 (The Edge ordering)**.** *The query complexity for* Crowd-ER *with the edge ordering algorithm is at most,*

$$n + \min_{1\leq s\leq n}\Big[\binom{k}{2}s^2 + n\sum_{i=1}^{k}\sum_{\ell=s}^{|V_i|}\ell L_{g,r}\Big(\binom{\ell}{2}\Big)\Big].$$

The proof of this theorem is provided in Section 5.

Theorem 3 (The Node ordering)**.** *The query complexity for* Crowd-ER *with the node ordering algorithm is at most,*

$$n + \sum_{i=1}^{k}\sum_{s=1}^{|V_i|}\min\{k, (n-|V_i|)L_{g,r}(s)\}.$$

The proof of this theorem is provided in Section 6.

4.1 Illustration: ϵ-biased Uniform Noise Model

We consider two distributions for f_r and f_g which are only ϵ far in terms of total variation distance from the uniform distribution. However, if we consider Hellinger distance, then **Dist-1** is closer to uniform distribution than **Dist-2**. These two distributions will be used as representative distributions to illustrate the potentials of the edge ordering and node ordering algorithms. In both cases, substituting ϵ with 0, we get uniform distribution which contains no information regarding the similarities of the entries.

Dist-1. Consider the following probability density functions for f_r and f_g, where $x \in [0,1]$, and $0 < \epsilon < 1/2$,

$$f_r(x) = \begin{cases}(1+\epsilon) & \text{if } x < \frac{1}{2}\\ (1-\epsilon) & \text{if } x \geq \frac{1}{2}\end{cases} \quad f_g(x) = \begin{cases}(1-\epsilon) & \text{if } x < \frac{1}{2}\\ (1+\epsilon) & \text{if } x \geq \frac{1}{2}.\end{cases}$$

Note that $\int_0^1 f_r(x)\,\mathrm{d}x = \int_0^{1/2}(1+\epsilon)\,\mathrm{d}x + \int_{1/2}^1(1-\epsilon)\,\mathrm{d}x = 1$. Similarly, $\int_0^1 f_g(x)\,\mathrm{d}x = 1$, that is they represent valid probability density functions. We have, $\mathcal{H}^2(f_g, f_r) = 1 - \int_0^1\sqrt{1-\epsilon^2}\mathrm{d}x = 1 - \sqrt{1-\epsilon^2} \approx \epsilon^2/2$.

Dist-2. Now consider the following probability density functions for f_r and f_g with $0 < \epsilon < 1/2$.

$$f_r(x) = \frac{1}{1-\epsilon}, 0 \leq x \leq 1-\epsilon, \quad f_g(x) = \frac{1}{1-\epsilon}, \epsilon \leq x \leq 1.$$

Again, $\int_0^1 f_r(x)\,\mathrm{d}x = \int_0^{1-\epsilon}\frac{1}{(1-\epsilon)}\,\mathrm{d}x = 1$. Similarly, $\int_0^1 f_g(x)\,\mathrm{d}x = \int_\epsilon^1 f_g(x)\,\mathrm{d}x = 1$, that is they represent valid probability density functions. We have, $\mathcal{H}^2(f_g, f_r) = 1 - \int_\epsilon^{1-\epsilon}\frac{1}{1-\epsilon}\mathrm{d}x = \frac{\epsilon}{1-\epsilon} \approx \epsilon$.

We have the following results for these two distributions.

Proposition 1 (Lower bound)**.** *Any (possibly randomized) algorithm for* Crowd-ER*, must make $\Omega(n + \frac{k^2}{\epsilon^2})$ queries for* **Dist-1** *and $\Omega(n + \frac{k^2}{\epsilon})$ queries for* **Dist-2***, to recover the clusters exactly (with probability 1).*

The proof of this theorem follows from Theorem 1, Corollary 1, and by plugging in the Hellinger distances between f_g, f_r in both cases.

The following set of results are corollaries of Theorem 2.

Proposition 2 (Uniform noise (no similarity information))**.** *Under the uniform noise model where $f_g, f_r \sim Unif[0,1]$, the edge ordering algorithm has query complexity $O(nk\log\frac{n}{k})$ for* Crowd-ER*.*

Proof. Since $f_g = f_r$, the similarity matrix W amounts to no information at all. We know that in this situation, one must make $O(nk)$ queries for the correct solution of Crowd-ER.

In this situation, a straight-forward calculation shows that, $L_{g,r}(t) = \frac{1}{t+1}$. This means, ignoring the first n term, from Theorem 2, the edge ordering algorithm makes at most $\min_{1\leq s\leq n}\Big[\binom{k}{2}s^2 + n\sum_{i=1}^k\sum_{\ell=s}^{|V_i|}\ell\frac{2}{\ell(\ell-1)+2}\Big] \leq \min_{1\leq s\leq n}\Big[\frac{k^2s^2}{2} + 2n\sum_{i=1}^k\sum_{\ell=s}^{|V_i|}\frac{1}{\ell-1}\Big]$ number of queries. By bounding the harmonic series and using the concavity of log, we have the number of queries made by the edge ordering algorithm is at most $\min_{1\leq s\leq n}\Big[\frac{k^2s^2}{2} + 2n\sum_{i=1}^k\ln\frac{|V_i|-1}{s-2}\Big]$ $\leq \min_{1\leq s\leq n}\Big[\frac{k^2s^2}{2} + 2nk\ln\frac{n-k}{k(s-2)}\Big] = O(nk\log\frac{n}{k})$, where we have substituted $s = \sqrt{n/k}$. □

Proposition 3 (Dist-1). *When $f_g, f_r \sim$ **Dist-1**, the edge ordering algorithm has query complexity $O(nk(1-2\epsilon)\log\frac{n}{k})$ for* Crowd-ER.

Proof. The proof is identical to the above. For small ϵ, we have $L_{g,r}(t) \approx \frac{1-\epsilon}{(1+\epsilon)(t+1)} \approx \frac{1-2\epsilon}{t+1}$ (see, Section 8). The algorithm queries at most $O(nk(1-2\epsilon)\log\frac{n}{k})$ edges. □

Proposition 4 (Dist-2). *When $f_g, f_r \sim$ **Dist-2**, the edge ordering algorithm has query complexity $O\left(n + \frac{k^2\log n}{\epsilon}\right)$ for* Crowd-ER.

Proof. For this case, we have $L_{g,r}(t) \leq \frac{e^{-\epsilon(t+1)}}{t+1}$ (see, Section 8). Choose $s = \sqrt{\frac{4\log n}{\epsilon}} + 1$. Then using Theorem 2, $n\sum_{i=1}^{k}\sum_{\ell=s}^{|V_i|}\ell L_{g,r}\left(\binom{\ell}{2}\right) \leq n\sum_{i=1}^{k}\sum_{\ell=s}^{|V_i|}\frac{2e^{-\epsilon\binom{\ell}{2}}}{\ell-1} < 1$. Therefore, the number of queries is $O\left(n + \frac{k^2\log n}{\epsilon}\right)$, matching the lower bound within a $\log n$ factor. □

For the Node-ordering algorithm, we have the following result as a corollary of Theorem 3.

Proposition 5 (Node-Ordering). *When $f_g, f_r \sim$ **Dist-1**, the node ordering algorithm has query complexity $O(nk(1-\epsilon^2))$ for* Crowd-ER. *When $f_g, f_r \sim$ **Dist-2**, node ordering has query complexity $O\left(n + \frac{k^2\log n}{\epsilon}\right)$ for* Crowd-ER.

Proof. For **Dist-1**, $L_{g,r}(s) \approx \frac{1-2\epsilon}{s+1}$. Therefore, when $s \geq \frac{n}{k}(1-\epsilon)$, $\min\{k, (n-|V_i|)L_{g,r}(s)\} \leq (1-\epsilon)k$. Thus, the total number of queries is $O(nk(1-\epsilon) + \epsilon nk(1-\epsilon)) = O(nk(1-\epsilon^2))$. For **Dist-2**, $L_{g,r}(s) = \frac{\exp(-\epsilon s)}{s+1}$. Therefore, when $s \geq \frac{2\log n}{\epsilon}$, $\min\{k, (n-|V_i|)L_{g,r}(s)\} \leq \frac{1}{n(s+1)}$. Thus the total number of expected queries is $O(n + \sum_{i=1}^{k}\frac{k\log n}{\epsilon} + \frac{|V_i|\log|V_i|}{n}) = O(n + \frac{k^2\log n}{\epsilon})$. □

Note that, there is no difference in the upper bounds given between the Edge and Node ordering algorithms for **Dist-2**. But Edge-ordering uses order $\log(n/k)$ factor more queries than the optimal ($O(nk)$) for **Dist-1**. **Dist-1** is closer to uniform distribution by the Hellinger measure than **Dist-2**, which shows that Hellinger distance is the right choice for distance here. Assuming $k = o(n)$, we get a drastic reduction in query complexity by moving from **Dist-1** to **Dist-2**.

5 Analysis of the Edge ordering algorithm: proof of Theorem 2

Let R be a random variable with distribution f_r and $G_1, \ldots, G_t$ be identical random variables with distribution f_g. Let $R, G_1, G_2, \ldots, G_t$ be all independent. Note that,

$$\Pr(R \geq \max\{G_1, \ldots, G_t\}) = \int_0^1 \Big(\int_0^r f_g(y)dy\Big)^t f_r(x)dx = L_{g,r}(t). \quad (1)$$

In the interest of clarity, let us call a pair $(u, v) \in V \times V$ a *green* edge iff $u, v \in V_i$ for some $i = 1, \ldots, k$, and otherwise call the pair a *red edge*.

In the current graph, let there exist ℓ nodes, called $U \subset V$, which all belong to the same cluster but no edge from the induced graph on these ℓ vertices have been queried yet. Then there are $\binom{\ell}{2}$ green edges within U, yet to be queried. On the other hand, there are at most $n\ell$ red edges with one end point incident on the vertices in U. We now count the number of red edges incident on U that the algorithm will query before querying a green edge within U. We can account for all the red edges queried by the algorithm by considering each cluster at a time, and summing over the queried red edges incident on it. In fact, by doing this, we double count every red edge. Since the probability of querying a red edge incident on U before querying any of the $\binom{\ell}{2}$ green edges incident on U is $L_{g,r}(\binom{\ell}{2})$, the expected number of queried red edges incident on U before querying a green edge in U is at most $n\ell L_{g,r}(\binom{\ell}{2}))$.

Let s be a positive integer. Consider a cluster $V_i : |V_i| \geq s$. Suppose at some point of time, there are ℓ components of V_i remaining to be connected. Then, again there are at least $\binom{\ell}{2}$ green edges, querying any of which will decrease the number of components by 1. Thus, the expected number of red edges that are queried incident on nodes in V_i before there remain at most s components of V_i is at most $n\sum_{\ell=s}^{|V_i|}\ell L_{g,r}(\binom{\ell}{2}))$. Therefore, the expected number of red edges that are queried until only s components are left for every cluster is $n\sum_{i=1}^{k}\sum_{\ell=s}^{|V_i|}\ell L_{g,r}\left(\binom{\ell}{2}\right)$.

Now the number of red edges across the clusters having size at most s is at most $\binom{k}{2}s^2$. Therefore, even if we query all those edges, we get the total number of queried red edges to be at most $\binom{k}{2}s^2 + n\sum_{i=1}^{k}\sum_{\ell=s}^{|V_i|}\ell L_{g,r}\left(\binom{\ell}{2}\right)$.

The algorithm queries a total of $n-k$ green edges, exactly spanning every cluster. Thus the total number of queries is at most $n + \binom{k}{2}s^2 + n\sum_{i=1}^{k}\sum_{\ell=s}^{|V_i|}\ell L_{g,r}\left(\binom{\ell}{2}\right)$.

6 Analysis of the Node ordering algorithm: proof of Theorem 3

The computed expected cluster size for each node can be a highly biased estimator, and may not provide any useful information. For example, the expected cluster size of a node in V_i is $\frac{\epsilon}{c}|V_i| + (\frac{1}{2} - \frac{\epsilon}{2c})n$ where $c = 2$ for **Dist 1** and $c = 1$ for **Dist 2**. Therefore, the node ordering considered by (Vesdapunt, Bellare, and Dalvi 2014) can be arbitrary. Hence, for the purpose of our analysis, we ignore this ordering based on the expected size.

Consider the state of the algorithm where it needs to insert a node v which truly belongs to cluster V_i. Suppose the current size of V_i is s, that is V_i already contains s nodes when v is considered. Consider another cluster V_j, $j \neq i$, and let its current size be s'. Let C_i and C_j denote the current subclusters of V_i and V_j that have been formed.Then, $P(w_{v,u} \geq \max_{x\in V_i} w_{v,x})$ where $u \in C_j$ is at most $L_{g,r}(s)$. Hence, $P(\exists u \in C_j, w_{v,u} \geq \max_{x\in V_i} w_{v,x}) \leq \min\{1, s'L_{g,r}(s)\}$. Thus the expected number of queried red edges before v is correctly inserted in V_i is at most $\min\{k, L_{g,r}(s)\sum_{j\in[1,k],j\neq i}|V_j|\} \leq$

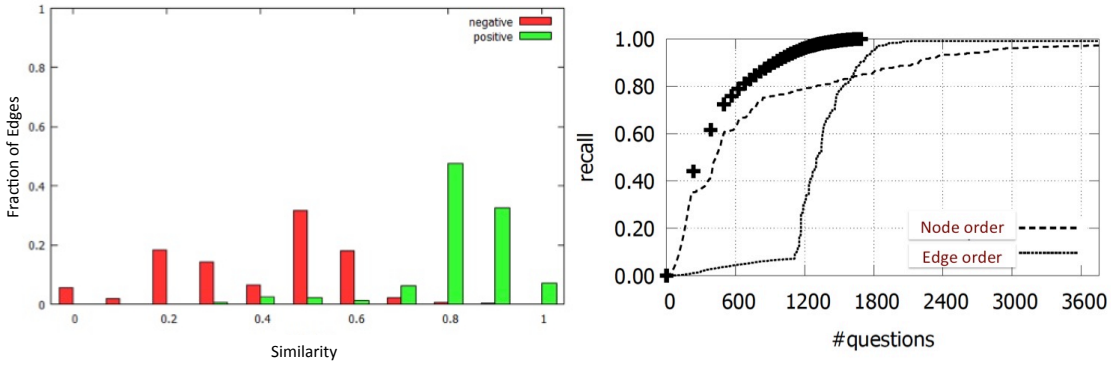

(a) `similarity value distribution` (b) `#queries vs recall`

Figure 2: `cora`

$\min\{k, (n - |V_i|)L_{g,r}(s)\}$. Hence the expected total number of queried red edges to grow the ith cluster is at most $\sum_{s=1}^{|V_i|} \min\{k, (n-|V_i|)L_{g,r}(s)\}$, and thus the expected total number of queries, including green and red edges is bounded by $n + \sum_{i=1}^{k} \sum_{s=1}^{|V_i|} \min\{k, (n - |V_i|)L_{g,r}(s)\}$.

7 Experimental Observations

A detailed comparison of the node ordering and edge ordering methods on multiple real datasets has been shown in (Vesdapunt, Bellare, and Dalvi 2014, Figures 12,14). The number of queries issued by the two methods are very close on complete resolution.To validate further, we did the following experiments.

Datasets. (i) We created multiple synthetic datasets each containing 1200 nodes and 14 clusters with the following size distribution: two clusters of size 200, four clusters of size 100, eight clusters of size 50, two clusters each of size 30 and 20 and the rest of the clusters of size 10. The datasets differed in the way similarity values are generated by varying ϵ and sampling the values either from **Dist-1** or **Dist-2**. The similarity values are further discretized to take values from the set $\{0, 0.1, 0.2, ..., 0.9, 1\}$.

(ii) We used the widely used `cora` (McCallum 2004) dataset for ER. `cora` is a bibliography dataset, where each record contains title, author, venue, date, and pages attributes. There are 1878 nodes in total with 191 clusters, among which 124 are non-singletons. The largest cluster size is 236, and the total number of pairs is $17,64,381$. We used the similarity function as in (Whang, Lofgren, and Garcia-Molina 2013; Wang et al. 2013; Vesdapunt, Bellare, and Dalvi 2014; Firmani, Saha, and Srivastava 2016).

Observation. The number of queries for the node-ordering and edge-ordering algorithms are reported in Table 1 for the synthetic datasets. Clearly, the number of queries asked for **Dist-2** is significantly less than that for **Dist-1** at the same value of ϵ. This confirms with our theoretical findings. Interestingly, we observe that the number of queries asked by the edge-ordering algorithm is consistently higher than the node-ordering algorithm under **Dist-1**. This is also expected from Propositions 3 and 5 due to a gap of $\log \frac{n}{k}$ in the number of queries of the two algorithms. In a similar vein, we see the edge-ordering algorithm is more effective than the node-ordering for **Dist-2**, possibly because of hidden constants in the asymptotic analysis.

Node-Ordering	Edge-Ordering	Distribution	ϵ
4475	4460	**Dist-1**	$\epsilon = \frac{1}{2}$
5207	6003	**Dist-1**	$\epsilon = \frac{1}{3}$
5883	7145	**Dist-1**	$\epsilon = \frac{1}{4}$
6121	7231	**Dist-1**	$\epsilon = \frac{1}{5}$
6879	8545	**Dist-1**	$\epsilon = \frac{1}{10}$
7398	9296	**Dist-1**	$\epsilon = \frac{1}{20}$
1506	1277	**Dist-2**	$\epsilon = \frac{1}{5}$
1986	1296	**Dist-2**	$\epsilon = \frac{1}{10}$
2760	1626	**Dist-2**	$\epsilon = \frac{1}{20}$

Table 1: Number of Queries for **Dist-1** and **Dist-2**

Figure 2(a) shows the similarity value distribution for `cora` which is closer to **Dist-2** than **Dist-1**. Figure 2(b) shows the recall vs number of queries issued by the two methods. The line marked with '+' sign is the curve for the ideal algorithm that will ask only the required "green" edges first to grow all the clusters and then ask just one "red" edge across every pair of clusters. Upon completion, the number of queries issued by the edge ordering and node ordering methods are respectively 21,099 and 23,243 which are very close to optimal. Interestingly, this confirms with our observation on the However, they achieve above 0.996 recall in less than $4,000$ queries. This can also be explained by our analysis. The remaining large number of queries are mainly spent on growing small clusters, e.g. when cluster sizes are $o(\log n)$– they do not give much benefit on recall, but consume many queries.

8 Appendix: $L_{g,r}(t)$ for Dist-1, Dist-2

Proposition 6. *For $f_g, f_r \sim$* **Dist-1** *and small ϵ, we have $L_{g,r}(t) \approx \frac{(1-\epsilon)}{(1+\epsilon)(t+1)}$.*

Proof. We have,

$$L_{g,r}(t) = \int_{r=0}^{1} \left(\int_{x=0}^{r} f_G(x)\,\mathrm{d}x \right)^t f_R(r)\,\mathrm{d}r$$

$$= \int_{r=0}^{1/2} \left(\int_{x=0}^{r} f_G(x)\,\mathrm{d}x \right)^t (1+\epsilon)\,\mathrm{d}r + \int_{r=1/2}^{1} \left(\int_{x=0}^{r} f_G(x)\,\mathrm{d}x \right)^t (1-\epsilon)\,\mathrm{d}r$$

$$= \int_{r=0}^{1/2} \left(\int_{x=0}^{r} (1-\epsilon)\,\mathrm{d}x \right)^t (1+\epsilon)\,\mathrm{d}r + \int_{r=1/2}^{1} \left(\int_{x=0}^{1/2} (1-\epsilon)\,\mathrm{d}x + \int_{x=1/2}^{r} (1+\epsilon)\,\mathrm{d}x \right)^t (1-\epsilon)\,\mathrm{d}r$$

$$= \frac{(1+\epsilon)(1-\epsilon)^t}{2^{t+1}(t+1)} + (1-\epsilon)\int_{r=1/2}^{1} (r(1+\epsilon)-\epsilon)^t\ \mathrm{d}r$$

Set $z = r(1+\epsilon)-\epsilon$, then $\mathrm{d}z = (1+\epsilon)\mathrm{d}r$. We have

$$(1-\epsilon)\int_{r=1/2}^{1} (r(1+\epsilon)-\epsilon)^t\ \mathrm{d}r = \frac{(1-\epsilon)}{(1+\epsilon)}\int_{z=\frac{(1-\epsilon)}{2}}^{1} z^t\,\mathrm{d}z = \frac{(1-\epsilon)}{(1+\epsilon)(t+1)}\left(1-\frac{(1-\epsilon)^{t+1}}{2^{t+1}}\right).$$

Therefore,

$$L_{g,r}(t) = \frac{(1+\epsilon)(1-\epsilon)^t}{2^{t+1}(t+1)} + \frac{(1-\epsilon)}{(1+\epsilon)(t+1)}\left(1-\frac{(1-\epsilon)^{t+1}}{2^{t+1}}\right) = \frac{(1-\epsilon)}{(1+\epsilon)(t+1)}\left(1+\epsilon\left(\frac{1-\epsilon}{2}\right)^{t-1}\right).$$

□

Proposition 7. *For $f_g, f_r \sim$* **Dist-2** *we have $L_{g,r}(t) \leq \frac{\mathrm{e}^{-\epsilon(t+1)}}{t+1}$.*

Proof. We have,

$$L_{g,r}(t) = \int_{r=\epsilon}^{1-\epsilon} \left(\int_{x=\epsilon}^{r} f_G(x)\,\mathrm{d}x \right)^t \frac{1}{1-\epsilon}\,\mathrm{d}r = \int_{r=\epsilon}^{1-\epsilon} \left(\int_{x=\epsilon}^{r} \frac{1}{1-\epsilon}\,\mathrm{d}x \right)^t \frac{1}{1-\epsilon}\,\mathrm{d}r$$

$$= \frac{1}{(1-\epsilon)^{t+1}}\int_{r=\epsilon}^{1-\epsilon} (r-\epsilon)^t\,\mathrm{d}r = \frac{1}{t+1}\left(\frac{1-2\epsilon}{1-\epsilon}\right)^{t+1} = \frac{1}{t+1}\left(1-\frac{\epsilon}{1-\epsilon}\right)^{t+1} \leq \frac{(1-\epsilon)^{t+1}}{t+1} \leq \frac{\mathrm{e}^{-\epsilon(t+1)}}{t+1}.$$

□

Acknowledgements: This research is supported in part by NSF CCF Awards 1464310, 1642658, 1642550 and a Google Research Award. The authors would like to thank Sainyam Galhotra for his many help with the simulation results.

References

[Bellare et al. 2012] Bellare, K.; Iyengar, S.; Parameswaran, A. G.; and Rastogi, V. 2012. Active sampling for entity matching. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '12, 1131–1139.

[Christen 2012] Christen, P. 2012. *Data matching: concepts and techniques for record linkage, entity resolution, and duplicate detection*. Springer Science & Business Media.

[Davidson et al. 2014] Davidson, S. B.; Khanna, S.; Milo, T.; and Roy, S. 2014. Top-k and clustering with noisy comparisons. *ACM Trans. Database Syst.* 39(4):35:1–35:39.

[Demartini, Difallah, and Cudré-Mauroux 2012] Demartini, G.; Difallah, D. E.; and Cudré-Mauroux, P. 2012. Zencrowd: leveraging probabilistic reasoning and crowdsourcing techniques for large-scale entity linking. In *WWW*, 469–478.

[Elmagarmid, Ipeirotis, and Verykios 2007] Elmagarmid, A. K.; Ipeirotis, P. G.; and Verykios, V. S. 2007. Duplicate record detection: A survey. *IEEE Trans. Knowl. Data Eng.* 19(1):1–16.

[Fellegi and Sunter 1969] Fellegi, I. P., and Sunter, A. B. 1969. A theory for record linkage. *Journal of the American Statistical Association* 64(328):1183–1210.

[Firmani, Saha, and Srivastava 2016] Firmani, D.; Saha, B.; and Srivastava, D. 2016. Online entity resolution using an oracle. *PVLDB* 9(5):384–395.

[Getoor and Machanavajjhala 2012] Getoor, L., and Machanavajjhala, A. 2012. Entity resolution: theory, practice & open challenges. *PVLDB* 5(12):2018–2019.

[Gruenheid et al. 2015] Gruenheid, A.; Nushi, B.; f, T.; Gatterbauer, W.; and Kossmann, D. 2015. Fault-tolerant entity resolution with the crowd. *CoRR* abs/1512.00537.

[Larsen and Rubin 2001] Larsen, M. D., and Rubin, D. B. 2001. Iterative automated record linkage using mixture models. *Journal of the American Statistical Association* 96(453):32–41.

[Mazumdar and Saha 2016] Mazumdar, A., and Saha, B. 2016. Clustering via crowdsourcing. *arXiv preprint arXiv:1604.01839*.

[McCallum 2004] McCallum, A. 2004. `https://people.cs.umass.edu/~mccallum/data/cora-refs.tar.gz`.

[Sarawagi and Bhamidipaty 2002] Sarawagi, S., and Bhamidipaty, A. 2002. Interactive deduplication using active learning. In *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '02, 269–278.

[Verroios and Garcia-Molina 2015] Verroios, V., and Garcia-Molina, H. 2015. Entity resolution with crowd errors. In *31st IEEE International Conference on Data Engineering, ICDE 2015, Seoul, South Korea, April 13-17, 2015*, 219–230.

[Vesdapunt, Bellare, and Dalvi 2014] Vesdapunt, N.; Bellare, K.; and Dalvi, N. 2014. Crowdsourcing algorithms for entity resolution. *PVLDB* 7(12):1071–1082.

[Wang et al. 2012] Wang, J.; Kraska, T.; Franklin, M. J.; and Feng, J. 2012. Crowder: Crowdsourcing entity resolution. *PVLDB* 5(11):1483–1494.

[Wang et al. 2013] Wang, J.; Li, G.; Kraska, T.; Franklin, M. J.; and Feng, J. 2013. Leveraging transitive relations for crowdsourced joins. In *SIGMOD Conference*, 229–240.

[Whang, Lofgren, and Garcia-Molina 2013] Whang, S. E.; Lofgren, P.; and Garcia-Molina, H. 2013. Question selection for crowd entity resolution. *PVLDB* 6(6):349–360.

[Yi et al. 2012] Yi, J.; Jin, R.; Jain, A. K.; Jain, S.; and Yang, T. 2012. Semi-crowdsourced clustering: Generalizing crowd labeling by robust distance metric learning. In *NIPS 2012.*, 1781–1789.